

>home > about > feedback > login

US Patent & Trademark Office



Try the new Portal design Give us your opinion after using it.

Search Results

Search Results for: [((buffer or queue) < near/8> (fixed or reserved)) and ((buffer or queue) <near/8>(pageable or virtual or logical)) and ((var* or alter* or chang* or adjust* or modif*) <near/6> (storage or memory or buffer or queue) <near/8> (compar* or equal* or different or similar))]

Found 31 of 111,550 searched.

Rerun within the Portal

Search within Results

> Advanced Search

> Search Help/Tips

Sort by: <u>Title Publication Publication Date</u>

Score

Binder

Results 1 - 20 of 31

short listing

1 Microprogramming revisited

99%

Michael J. Flynn, M. Donald McLaren

Proceedings of the 1967 22nd national conference January 1967

It is the objective of this paper to briefly trace the history of the idea and the difficulties involved with defining or implementing it. In doing this, we first consider the general control problem and instruction formats. Next, storage implementations of the control function are considered and a restricted definition of microprogramming is proposed. This is then evaluated from a technological, architectural and programming point of view. We hope to show that our (demanding) definition of ...

Synchronization and communication in the T3E multiprocessor

99%

বী Steven L. Scott

Proceedings of the seventh international conference on Architectural support for programming languages and operating systems September 1996

Volume 31, 30 Issue 9, 5

This paper describes the synchronization and communication primitives of the Cray T3E multiprocessor, a shared memory system scalable to 2048 processors. We discuss what we have learned from the T3D project (the predecessor to the T3E) and the rationale behind changes made for the T3E. We include performance measurements for various aspects of communication and synchronization. The T3E augments the memory interface of the DEC 21164 microprocessor with a large set of explicitly-managed, external r ...

3 Dynamic queue length thresholds for multiple loss priorities

Ellen L. Hahne, Abhijit K. Choudhury

98%



IEEE/ACM Transactions on Networking (TON) June 2002

Volume 10 Issue 3

Buffer management schemes are needed in shared-memory packet switches to regulate the sharing of memory among different output port queues and among traffic classes with different loss priorities. Earlier we proposed a single-priority scheme called Dynamic Threshold (DT), in which the maximum permissible queue length is proportional to the unused buffering in the switch. A queue whose length equals or exceeds the current threshold value may accept no new arrivals. In this paper, we propose, anal ...

4 Interconnect Architecture: Circuit design of routing switches

98%

Guy Lemieux, David Lewis

Proceedings of the 2002 ACM/SIGDA tenth international symposium on Field-programmable gate arrays February 2002

This paper examines circuit design of buffered routing switches in symmetrical, island-style FPGAs. The effects of switch size, tile length, level-restoring, and slow input slew rates are examined. Two new fanin-based switch designs are used to eliminate nearly all of the increase in delay that arises from fanout with a previous switch design. Alternating between buffers and pass transistors is shown to improve connection delay without fanout by 25%. To take advantage of this, we propose schemes ...

5 Effective graph clustering for path queries in digital map databases

98%

Yun-Wu Huang, Ning Jing, Elke A. Rundensteiner

Proceedings of the fifth international conference on Information and knowledge management November 1996

6 A decomposition-based simulated annealing technique for data clustering

98%

d Kien A. Hua, S. D. Lang, Wen K. Lee

Proceedings of the thirteenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems May 1994

It has been demonstrated that simulated annealing provides high-quality results for the data clustering problem. However, existing simulated annealing schemes are memory-based algorithms; they are not suited for solving large problems such as data clustering which typically are too big to fit in the memory space in its entirety. Various buffer replacement policies, assuming either temporal or spatial locality, are not useful in this case since simulated annealing is based o ...

7 Efficient implementation of event sets in Time Warp

98%

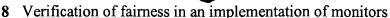
Robert Rönngren, Rassul Ayani, Richard M. Fujimoto, Samir R. Das

ACM SIGSIM Simulation Digest , Proceedings of the seventh workshop on Parallel and distributed simulation $\rm July~1993$

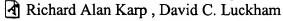
Volume 23 Issue 1

The implementation of the pending event set (PES) is crucial to the execution speed of discrete event simulation programs. This paper studies the implementation of the PES in the context of simulations executing on parallel computers using the Time Warp mechanism. We present a scheme for implementing Time Warsp's PES based on well-known data structures for priority queues. This scheme supports efficient management of future and past events, especially for rollback and fossil collection oper ...





97%



Proceedings of the 2nd international conference on Software engineering October 1976 An implementation in Pascal by Saxena and Bredt of the Monitor construct is studied. Techniques are given for using a program verifer to analyse the conditions under which the implementation is fair (i.e. once a process is delayed it eventually will be continued). By use of a virtual data structure, fairness is represented in terms of simple properties which can be verified automatically. Examples are given illustrating how the verification can force unstated assumptions upon whi ...

9 Parametric query optimization

97%

Yannis E. Ioannidis, Raymond T. Ng, Kyuseok Shim, Timos K. Sellis
The VLDB Journal — The International Journal on Very Large Data Bases May
1997

Volume 6 Issue 2

In most database systems, the values of many important run-time parameters of the system, the data, or the query are unknown at query optimization time. Parametric query optimization attempts to identify at compile time several execution plans, each one of which is optimal for a subset of all possible values of the run-time parameters. The goal is that at run time, when the actual parameter values are known, the appropriate plan should be identifiable with essentially no overhead. We present a g ...

10 Pushout with virtual thresholds buffer management scheme in a shared buffer ATM switch

, 97%

Ruey-Bin Yang, Yuan-Sun Chu, Cheng-Shong Wu, Ming-Cheng Liang International Journal of Network Management March 2003

Volume 13 Issue 2

In this paper we investigate the performance metrics of buffer management schemes. In general, the selective pushout (SP) scheme can support very low loss probability of the high-priority cells, but it may cause unfairness of buffer allocation among different output queues and high overall cell loss probability. In order to fit the dynamic required performance metrics of ATM switches, a novel buffer management scheme called pushout with virtual thresholds (PVT) is proposed. In the PVT scheme, ea ...

11 Dynamic speculation and synchronization of data dependences

97%

Andreas Moshovos, Scott E. Breach, T. N. Vijaykumar, Gurindar S. Sohi
ACM SIGARCH Computer Architecture News, Proceedings of the 24th annual international symposium on Computer architecture May 1997

Volume 25 Issue 2

Data dependence speculation is used in instruction-level parallel (ILP) processors to allow early execution of an instruction before a logically preceding instruction on which it may be data dependent. If the instruction is independent, data dependence speculation succeeds; if not, it fails, and the two instructions must be synchronized. The modern dynamically scheduled processors that use data dependence speculation do so blindly (i.e., every load instruction with unresolved dependences is spec ...

12 Performance analysis of recovery techniques

Andreas Reuter

ACM Transactions on Database Systems (TODS) December 1984

97%



Various logging and recovery techniques for centralized transaction-oriented database systems under performance aspects are described and discussed. The classification of functional principles that has been developed in a companion paper is used as a terminological basis. In the main sections, a set of analytic models is introduced and evaluated in order to compare the performance characteristics of nine different recovery techniques with respect to four key parameters and a set of other pa ...

13 Streaming 2: Server-based smoothing of variable bit-rate streams

97%

Stergios V. Anastasiadis, Kenneth C. Sevcik, Michael Stumm

Proceedings of the ninth ACM international conference on Multimedia October 2001 We introduce an algorithm that uses buffer space available at the server for smoothing disk transfers of variable bit-rate streams. Previous smoothing techniques prefetched stream data into the client buffer space, instead. However, emergence of personal computing devices with widely different hardware configurations means that we should not always assume abundance of resources at the client side. The new algorithm is shown to have optimal smoothing effect under the specified constraints. We inc ...

14 Pipeline Architecture

97%

C. V. Ramamoorthy, H. F. Li

ACM Computing Surveys (CSUR) January 1977

Volume 9 Issue 1

15 A buffer allocation scheme for ATM networks: complete sharing based on virtual partition

97%

Guo-Liang Wu, Jon W. Mark

IEEE/ACM Transactions on Networking (TON) December 1995

Volume 3 Issue 6

16 Segmented FIFO page replacement

96%

Rollins Turner, Henry Levy

Proceedings of the 1981 ACM SIGMETRICS conference on Measurement and modeling of computer systems September 1981

A fixed-space page replacement algorithm is presented. A variant of FIFO management using a secondary FIFO buffer, this algorithm provides a family of performance curves lying between FIFO and LRU. The implementation is simple, requires no periodic scanning, and uses no special hardware support. Simulations are used to determine the performance of the algorithm for several memory reference traces. Both the fault rates and overhead cost are examined.

17 Join algorithm costs revisited

96%

Evan P. Harris, Kotagiri Ramamohanarao

The VLDB Journal — The International Journal on Very Large Data Bases January 1996

Volume 5 Issue 1

A method of analysing join algorithms based upon the time required to access, transfer and perform the relevant CPU-based operations on a disk page is proposed. The costs of variations of several of the standard join algorithms, including nested block, sort-merge, GRACE hash and hybrid hash, are presented. For a given total buffer size, the cost of these join algorithms depends on the parts of the buffer allocated for each purpose. For example, when joining two



relations using the nested block j ...

18 Design and analysis of low-power access protocols for wireless and mobile ATM networks

96%

Krishna M. Sivalingam, Jyh-Cheng Chen, Prathima Agrawal, Mani B. Srivastava Wireless Networks January 2000

Volume 6 Issue 1

This paper describes the design and analysis of a low‐ power medium access control (MAC) protocol for wireless/mobile ATM networks. The protocol – denoted EC‐MAC (energy conserving medium access control) – is designed to support different traffic types with quality‐of‐service (QoS) provisions. The network is based on the infrastructure model where a base station (BS) serves all the mobiles currently in its cell. A reserv ...

19 CARAT: a testbed for the performance evaluation of distributed database systems

96%

Walt Kohler, Bao-Chyuan Jenq

Proceedings of 1986 fall joint computer conference on Fall joint computer conference November 1999

20 Prefetching in a texture cache architecture

96%

Homan Igehy, Matthew Eldridge, Kekoa Proudfoot

Proceedings of the 1998 EUROGRAPHICS/SIGGRAPH workshop on Graphics hardware August 1998

Results 1 - 20 of 31

short listing



The ACM Portal is published by the Association for Computing Machinery. Copyright © 2003 ACM, Inc.



US Patent & Trademark Office



Try the new Portal design

Give us your opinion after using it.

Search Results

Search Results for: [((buffer or queue) <near/8> (fixed or reserved)) and ((buffer or queue) <near/8>(pageable or virtual or logical)) and ((var* or alter* or chang* or adjust* or modif*) <near/6> (storage or memory or buffer or queue) <near/8> (compar* or equal* or different or similar))]

Found 31 of 111,550 searched. → Rerun within the Portal

Search within Results

> Advanced Search

> Search Help/Tips

Sort by:

Title Publication Publication Date

Score

Binder

Results 21 - 31 of 31 short listing



21 Approximate time-parallel simulation of queueing systems with losses

96%

Jain J. Wang, Marc Abrams

Proceedings of the 24th conference on Winter simulation December 1992

22 The DARPA wideband network protocol

95%

W. Edmond, K. Leib, C. Topolcic

ACM SIGCOMM Computer Communication Review, Proceedings of the ACM symposium on Communications architectures & protocols August 1990

Volume 20 Issue 4

This paper discusses the Dual Bus Protocol (DBP) developed for the Terrestrial Wideband Network (TWBNET), a high-speed trans-continental network that supports distributed real-time applications and protocol research. DBP is a descendant of the QPSX protocol proposed as the IEEE 802.6 Metropolitan Area Network (MAN) standard, with enhancements to create a protocol suitable for a wide area network and to provide network services such as reserved bandwidth. In addition, the protocol was design ...

23 Design of transaction-oriented systems employing a transaction monitor

95%

R. A. Davenport

Proceedings of the 1974 annual conference January 1974

An increasing number of computer users are implementing an on-line facility with the aid of a transaction monitor program. This transaction monitor may be supplied either by the computer hardware manufacturer or by a software supplier. It is felt that the literature is inadequate in aiding such a user to implement his system. Consequently this paper sets out to achieve such an aim by presenting a methodology for the design of a transaction-oriented system employing a

transaction monitor. Th ...

24 Combining produce and consume operations in a pipelined shared memory multiprocessor

95%

B. J. Rodriguez, H. F. Jordan

Proceedings of the 1989 ACM/IEEE conference on Supercomputing August 1989 A technique for combining produce and consume operations is proposed. An Omega network was designed to work with a pipelined micro multiprocessor system, and each of its nodes to be implemented on a single chip. Without a significant increase in hardware complexity, the nodes combine produce and consume messages in an unbounded way. The combining technique that was developed guarantees that every produce and consume message that is combined will always succeed, adding a new benefit to combi ...

25 Buffer management based on return on consumption in a multi-query environment

95%

Philip S. Yu, Douglas W. Cornell

The VLDB Journal — The International Journal on Very Large Data Bases January 1993

Volume 2 Issue 1

In a multi-query environment, the marginal utilities of allocating additional buffer to the various queries can be vastly different. The conventional approach examines each query in isolation to determine the optimal access plan and the corresponding locality set. This can lead to performance that is far from optimal. As each query can have different access plans with dissimilar locality sets and sensitivities to memory requirement, we employ the concepts of memory consumption and return on cons ...

26 Detecting graph-based spatial outliers: algorithms and applications (a summary of results)

95%

Shashi Shekhar, Chang-Tien Lu, Pusheng Zhang

Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining August 2001

Identification of outliers can lead to the discovery of unexpected, interesting, and useful knowledge. Existing methods are designed for detecting spatial outliers in multidimensional geometric data sets, where a distance metric is available. In this paper, we focus on detecting spatial outliers in graph structured data sets. We define statistical tests, analyze the statistical foundation underlying our approach, design several fast algorithms to detect spatial outliers. and provide a cost model ...

27 Dynamic-CBT and ChIPS— router support for improved multimedia performance on the 95%

1 Internet

Jae Chung, Mark Claypool

Proceedings of the eighth ACM international conference on Multimedia October 2000

The explosive increase in the volume and variety of Internet traffic has placed a growing emphasis on congestion control and fairness in Internet routers. Approaches to the problem of congestion, such as active queue management schemes like Random Early Detection (RED) use congestion avoidance techniques and are successful with TCP flows. Approaches to the problem of fairness, such as Fair Random Early Drop (FRED), keep per-flow state and punish misbehaved, non-TCP flows. Unfortunately, these ...



management in cell-switching networks

Allen R. Bonde, Sumit Ghosh

IEEE/ACM Transactions on Networking (TON) August 1994

Volume 2 Issue 4

29 High-speed switch scheduling for local-area networks

95%

Thomas E. Anderson, Susan S. Owicki, James B. Saxe, Charles P. Thacker

ACM Transactions on Computer Systems (TOCS) November 1993

Volume 11 Issue 4

Current technology trends make it possible to build communication networks that can support high-performance distributed computing. This paper describes issues in the design of a prototype switch for an arbitrary topology point-to-point network with link speeds of up to 1 Gbit/s. The switch deals in fixed-length ATM-style cells, which it can process at a rate of 37 million cells per second. It provides high bandwidth and low latency for datagram traffic. In addition, it supports real-time t ...

30 Query evaluation techniques for large databases

95%

4 Goetz Graefe

ACM Computing Surveys (CSUR) June 1993

Volume 25 Issue 2

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

31 Empirical performance evaluation of concurrency and coherency control protocols for database 95%

sharing systems

Erhard Rahm

ACM Transactions on Database Systems (TODS) June 1993

Volume 18 Issue 2

Database Sharing (DB-sharing) refers to a general approach for building a distributed high performance transaction system. The nodes of a DB-sharing system are locally coupled via a high-speed interconnect and share a common database at the disk level. This is also known as a " shared disk" approach. We compare database sharing with the database partitioning (shared nothing) approach and discuss the functional DBMS components that require new and coordinated solutions for DB-shar ...

Results 21 - 31 of 31

short listing



The ACM Portal is published by the Association for Computing Machinery. Copyright © 2003 ACM, Inc.